

Vector-Space Markov Random Fields via Exponential Families

Wesley Tansey*
Oscar Hernan Madrid Padilla†
Arun Sai Suggala‡
Pradeep Ravikumar§

Abstract

We present Vector-Space Markov Random Fields (VS-MRFs), a novel class of undirected graphical models where each variable can belong to an arbitrary vector space. VS-MRFs generalize a recent line of work on scalar-valued, uni-parameter exponential family and mixed graphical models, thereby greatly broadening the class of exponential families available (e.g., allowing multinomial and Dirichlet distributions). Specifically, VS-MRFs are the joint graphical model distributions where the node-conditional distributions belong to generic exponential families with general vector space domains. We also present a sparsistent M -estimator for learning our class of MRFs that recovers the correct set of edges with high probability. We validate our approach via a set of synthetic data experiments as well as a real-world case study of over four million foods from the popular diet tracking app MyFitnessPal. Our results demonstrate that our algorithm performs well empirically and that VS-MRFs are capable of capturing and highlighting interesting structure in complex, real-world data. All code for our algorithm is open source and publicly available.

1 Introduction

Undirected graphical models, also known as Markov Random Fields (MRFs), are a popular class of models for probability distributions over random vectors. Popular parametric instances include Gaussian MRFs, Ising, and Potts models, but these are all suited to specific data-types: Ising models for binary data, Gaussian MRFs for thin-tailed continuous data, and so on. Conversely, when there is prior knowledge of the graph structure but limited information otherwise, nonparametric approaches are available Sudderth et al. [2010]. A recent line of work has considered the challenge of specifying classes of MRFs targeted to the data-types in the given application, when

*Department of Computer Science, University of Texas at Austin, tansey@cs.utexas.edu (corresponding author)

†Department of Statistics and Data Sciences; University of Texas at Austin, oscar.madrid@utexas.edu

‡Department of Computer Science, University of Texas at Austin, arunsai@utexas.edu

§Department of Computer Science, University of Texas at Austin, pradeepr@cs.utexas.edu

the structure is unknown. For the specific case of homogeneous data, where each variable in the random vector has the same data-type, Yang et al. [2012] proposed a general subclass of homogeneous MRFs. In their construction, they imposed the restriction that each variable conditioned on other variables belong to a shared exponential family distribution, and then performed a Hammersley-Clifford-like analysis to derive the corresponding joint graphical model distribution, consistent with these node-conditional distributions. As they showed, even classical instances belong to this sub-class of MRFs; for instance, with Gaussian MRFs and Ising models, the node-conditional distributions follow univariate Gaussian and Bernoulli distributions respectively.

Yang et al. [2014] then proposed a class of *mixed MRFs* that extended this construction to allow for random vectors with variables belonging to different data types, and allowing each node-conditional distribution to be drawn from a different univariate, uni-parameter exponential family member (such as a Gaussian with known variance or a Bernoulli distribution). This flexibility in allowing for different univariate exponential family distributions yielded a class of mixed MRFs over heterogeneous random vectors that were capable of modeling a much wider class of distributions than was previously feasible, opening up an entirely new suite of possible applications.

To summarize, the state of the art can specify MRFs over heterogeneous data-typed random vectors, under the restriction that each variable conditioned on others belong to a uni-parameter, univariate exponential family distribution. But in many applications, such a restriction would be too onerous. For instance, a discrete random variable is best modeled by a categorical distribution, but this is a *multi-parameter* exponential family distribution, and does not satisfy the required restriction above. Other multi-parameter exponential family distributions popular in machine learning include gamma distributions with unknown shape parameter and Gaussian distributions with unknown variance. Another restriction above is that the variables be scalar-valued; but in many applications the random variables could belong to more general vector spaces, for example a Dirichlet distribution.

As modern data modeling requirements evolve, extending MRFs beyond such restrictive paradigms is becoming increasingly important. In this paper, we thus extend the above line of work in Yang et al. [2012, 2014]. As opposed to other approaches which merely cluster scalar variables Vats and Moura [2012], we allow node-conditional distributions to belong to a generic exponential family with a general vector space domain. We then perform a subtler Hammersley-Clifford-like analysis to derive a novel class of vector-space MRFs (VS-MRFs) as joint distributions consistent with these node-conditional distributions. This class of VS-MRFs provides support for the many modelling requirements outlined above, and could thus greatly expand the potential applicability of MRFs to new scientific analyses.

We also introduce an M -estimator for learning this class of VS-MRFs based on the sparse group lasso, and show that it is sparsistent, and that it succeeds in recovering the underlying edges of the graphical model. To solve the M -estimation problem, we also provide a scalable optimization algorithm based on Alternating Direction Method of Multipliers (ADMM) Boyd et al. [2011]. We validate our approach empirically via synthetic experiments measuring performance across a variety of scenarios. We also demonstrate the usefulness of VS-MRFs by modeling a real-world dataset of over four million foods from the MyFitnessPal food database.

The remainder of this paper is organized as follows. Section 2 provides background on mixed MRFs in the uni-parameter, univariate case. Section 3 details our generalization of the mixed MRF derivations to the vector-space case. Section 4 introduces our M -estimator and derives its sparsistency statistical guarantees. Section 5 contains our synthetic experiments and the MyFitnessPal case study. Finally, Section 6 presents concluding remarks and potential future work.

2 Background: Scalar Mixed Graphical Models

Let $X = (X_1, X_2, \dots, X_p)$ be a p -dimensional random vector, where each variable X_r has domain \mathcal{X}_r . An undirected graphical model or a Markov Random Field (MRF) is a family of joint distributions over the random vector X that is specified by a graph $G = (V, E)$, with nodes corresponding to each of the p random variables $\{X_r\}_{r=1}^p$, and edges that specify the factorization of the joint as:

$$\mathbb{P}(X) \propto \prod_{C \in \mathcal{C}(G)} \psi_C(X_C),$$

where $\mathcal{C}(G)$ is the set of fully connected subgraphs (or cliques) of the graph G , $X_C = \{X_s\}_{s \in C}$ denotes the subset of variables in the subset $C \subseteq V$, and $\{\psi_C(X_C)\}_{C \in \mathcal{C}(G)}$ are *clique-wise* functions, each of which is a “local function” in that it only depend on the variables in the corresponding clique, so that $\psi_C(X_C)$ only depends on the variable subset X_C .

Gaussian MRFs, Ising MRFs, etc. make particular parametric assumptions on these clique-wise functions, but a key question is whether there exists a more flexible specification of the form of these clique-wise functions that is targeted to the data-type and other characteristics of the random vector X .

For the specific case where the variables are scalars, so that the domains $\mathcal{X}_r \subseteq \mathbb{R}$, in a line of work, Yang et al. [2012, 2014] used the following construction to derive a subclass of MRFs targeted to the random vector X . Suppose that for variables $X_r \in \mathcal{X}_r$, the following (single-parameter) univariate exponential family distribution $P(X_r) = \exp\{\theta_r B_r(X_r) + C_r(X_r) - A_r(\theta_r)\}$, with natural parameter scalar θ , sufficient statistic scalar $B_r(X_r)$, base measure $C_r(X_r)$ and log normalization constant $A_r(\theta)$, serves as a suitable statistical model. Suppose that we use these univariate distributions to specify *conditional* distributions:

$$P(X_r | X_{-r}) = \exp\left\{ \begin{array}{l} E_r(X_{-r})B_r(X_r) + \\ C_r(X_r) - A_r(X_{-r}) \end{array} \right\} \quad , \quad (1)$$

where $E_r(\cdot)$ is an arbitrary function of the rest of the variables X_{-r} that serves as the natural parameter. Would these node-conditional distributions for each node $r \in V$ be consistent with some joint distribution for some specification of these functions $\{E_r(\cdot)\}_{r \in V}$? Theorem 1 from Yang et al. [2014] shows that there does exist a unique

joint MRF distribution with the form:

$$\begin{aligned}
P(X; \theta) = \exp \bigg\{ & \sum_{r \in V} \theta_r B_r(X_r) \\
& + \sum_{r \in V} \sum_{t \in N(r)} \theta_{rt} B_t(X_t) B_r(X_r) + \dots \\
& + \sum_{(t_1, \dots, t_k) \in \mathcal{C}} \theta_{t_1 \dots t_k}(X) \prod_{j=1}^k B_{t_j}(X_{t_j}) \\
& + \sum_{r \in V} C_r(X_r) - A(\theta) \bigg\}, \tag{2}
\end{aligned}$$

where $A(\theta)$ is the log-normalization constant. Their proof followed an analysis similar to the Hammersley-Clifford Theorem Lauritzen [1996], and entailed showing that for a consistent joint, the only feasible conditional parameter functions $E_r(\cdot)$ had the following form:

$$\begin{aligned}
E_r(X_{-r}) = & \theta_r + \sum_{t \in N(r)} \theta_{rt} B_t(X_t) + \dots \\
& + \sum_{t_2, \dots, t_k \in N(r)} \theta_{rt_2 \dots t_k}(X) \prod_{j=2}^k B_{t_j}(X_{t_j}), \tag{3}
\end{aligned}$$

where $\theta_{r\cdot} := \{\theta_r, \theta_{rt}, \dots, \theta_{rt_2 \dots t_k}\}$ is a set of parameters, and $N(r)$ is the set of neighbors of node r .

While their construction allows the specification of targeted classes of graphical models for heterogeneous random vectors, the conditional distribution of each variable conditioned on the rest of the variables is assumed to be a single-parameter exponential family distribution with a scalar sufficient statistic and natural parameter. Furthermore, their Hammersley-Clifford type analysis and sparsistency proofs relied crucially on that assumption. However in the case of multi-parameter and multivariate distributions, the sufficient statistics are a vector; indeed the random variables need not be scalars at all but could belong to a more general vector space. Could one construct classes of MRFs for this more general, but prevalent, setting? In the next section, we answer in the affirmative, and present a generalization of mixed MRFs to the vector-space case, with support for more general exponential families.

3 Generalization to the Vector-space Case

Let $X = (X_1, X_2, \dots, X_p)$ be a p -dimensional random vector, where each variable X_r belongs to a vector space \mathcal{X}_r . As in the scalar case, we will assume that a suitable statistical model for variables $X_r \in \mathcal{X}_r$ is an exponential family distribution

$$P(X_r) = \exp \left\{ \sum_{j=1}^{m_r} \theta_{rj} B_{rj}(X_r) + C_r(X_r) - A_r(\theta) \right\}, \tag{4}$$

with natural parameters $\{\theta_{rj}\}_{j=1}^{m_r}$, and sufficient statistics $\{B_{rj}\}_{j=1}^{m_r}$, base measure $C_r(X_r)$ and log normalization constant $A_r(\theta)$. We assume the sufficient statistics $B_{rj} : \mathcal{X}_r \mapsto \mathbb{R}$ lie in some Hilbert space \mathcal{H}_s , and moreover specify a minimal exponential family so that:

$$\sum_{j=1}^{m_r} \alpha_j B_{rj}(X_r) \neq c, \quad (5)$$

for any constant c and any vector $\alpha \neq \mathbf{0}$. We note that even though the variables $\{X_r\}$ could lie in general vector spaces, the exponential family distribution above is finite-dimensional. However, it has multiple parameters, which is the other facet that distinguishes it from the single-parameter univariate setting of Yang et al. [2012, 2014]. We defer a generalization of our framework to infinite-dimensional exponential families to future work.

Suppose we use these general exponential family distributions to specify node-conditional distributions of variables X_r conditioned on the rest of the random variables:

$$P(X_r|X_{-r}) = \exp\left\{\sum_{j=1}^{m_r} E_{rj}(X_{-r})B_{rj}(X_r) + C_r(X_r) - A_r(X_{-r})\right\}, \quad (6)$$

where $\{E_{rj}(X_{-r})\}_{j=1}^{m_r}$ are arbitrary functions of the rest of the variables that serve as natural parameters for the conditional distribution of X_r . As before, we ask the question whether these node-conditional distributions can be consistent with some joint distribution for some specification of the parameter functions $\{E_{rj}(X_{-r})\}_{j=1}^{m_r}$; the following theorem addresses this very question.

Theorem 1. *Let $X = (X_1, X_2, \dots, X_p)$ be a p -dimensional random vector with node-conditional distribution of each random vector X_r conditioned on the rest of random variables as defined in (6). These node-conditionals are consistent with a joint MRF distribution over the random vector X , that is, Markov with respect to a graph $G = (V, E)$ with clique-set \mathcal{C} , and with factors of size at most k , **if and only if** the functions $\{E_r(\cdot)\}_{r \in V}$ specifying the node-conditional distributions have the form:*

$$E_{ri}(X_{-r}) = \theta_{ri} + \sum_{t \in N(r)} \sum_{j=1}^{m_t} \theta_{ri;tj} B_{tj}(X_t) + \dots + \sum_{\substack{t_2, \dots, t_k \in N(r) \\ i_2=1 \dots m_{t_2} \\ \vdots \\ i_k=1 \dots m_{t_k}}} \theta_{ri; \dots; t_k i_k} \prod_{j=2}^k B_{t_j i_j}(X_{t_j}), \quad (7)$$

where $\theta_r = \{\theta_{ri}, \theta_{ri;tj}, \theta_{ri; \dots; t_k i_k}\}$ is a set of parameters, m_t is the dimension of the sufficient statistic vector for the t^{th} node-conditional distribution, and $N(r)$ is the set of neighbors of node r in graph G . The corresponding consistent joint MRF distribution has the following form:

$$\begin{aligned}
P(X|\theta) = \exp \Bigg\{ & \sum_{r \in V} \sum_{i=1}^{m_r} \theta_{ri} B_{ri}(X_r) + \dots \\
& + \sum_{t_1, \dots, t_k \in C} \sum_{\substack{i_1=1 \dots m_{t_1} \\ \vdots \\ i_k=1 \dots m_{t_k}}} \theta_{t_1 i_1, \dots, t_k i_k} \prod_{j=1}^k B_{t_j i_j}(X_{t_j}) \\
& + \sum_{r \in V} C_r(X_r) - A(\theta) \Bigg\} . \tag{8}
\end{aligned}$$

We provide a Hammersley-Clifford type analysis as proof of this theorem in the supplementary material, which however has subtleties not present in Yang et al. [2012, 2014], due to the arbitrary vector space domain of X_r , and the multiple parameters in the exponential families, which consequently entailed leveraging the geometry of the corresponding Hilbert spaces $\{\mathcal{H}_s\}_{s \in V}$ underlying the sufficient statistics $\{B_{sj}\}$.

The above Theorem 1 provides us with a general class of vector-space MRFs (VS-MRFs), where each variable could belong to more general vector space domains, and whose conditional distributions are specified by more general finite-dimensional exponential families. Consequently, many common distributions can be incorporated into VS-MRFs that were previously unsupported or lacking in Yang et al. [2012, 2014]. For instance, gamma and Gaussian nodes, though univariate, require vector-space parameters in order to be fully modeled. Additionally, multivariate distributions that were impossible to use with previous methods, such as the multinomial and Dirichlet distributions are now also available.

3.1 Pairwise conditional and joint distributions

Given the form of natural parameters in (7), the conditional distribution of a node X_r given all other nodes X_{-r} for the special case of pairwise MRFs (i.e. $k = 2$) has the form

$$\begin{aligned}
P(X_r|X_{-r}, \theta_r, \theta_{rt}) &= \exp \left\{ \sum_{i=1}^{m_r} \theta_{ri} B_{ri}(X_r) \right. \\
&+ \sum_{t \in N(r)} \sum_{i=1}^{m_r} \sum_{j=1}^{m_t} \theta_{ri;tj} B_{tj}(X_t) B_{ri}(X_r) \\
&\left. + C_r(X_r) - A_r(X_{-r}, \theta_r) \right\} , \tag{9} \\
&= \exp \left\{ \left\langle B_r(X_r), \theta_r + \sum_{t \in N(r)} \theta_{rt} B_t(X_t) \right\rangle \right. \\
&\left. + C_r(X_r) - A_r \left(\theta_r + \sum_{t \in N(r)} \theta_{rt} B_t(X_t) \right) \right\}
\end{aligned}$$

where θ_r is a vector formed from scalars $\{\theta_{ri}\}_{i=1}^{m_r}$, θ_{rt} is a matrix of dimension $m_r \times m_t$ obtained from scalars $\theta_{ri;tj}$ and $\langle \cdot, \cdot \rangle$ represents dot product between two vectors. Thus,

the joint distribution has the form

$$P(X|\theta) = \exp \left\{ \sum_{r \in V} \left\langle B_r(X_r), \theta_r + \sum_{t \in N(r)} \theta_{rt} B_t(X_t) \right\rangle + \sum_{r \in V} C_r(X_r) - A(\theta) \right\}, \quad (10)$$

with the log-normalization constant $A(\theta) = \log \int_{\mathcal{X}} \exp \{ \sum_{r \in V} \langle B_r(X_r), \theta_r + \sum_{t \in N(r)} \theta_{rt} B_t(X_t) \rangle + \sum_{r \in V} C_r(X_r) \}$. Since $A(\theta)$ is generally intractable to calculate, we next present an efficient approach to learning the structure of VS-MRFs.

4 Learning VS-MRFs

To avoid calculation of the log-normalization constant, we approximate the joint distribution in (10) with the independent product of node conditionals, also known as the *pseudo-likelihood*,

$$P(X|\theta) \approx \prod_r P(X_r | X_{-r}, \theta_r, \theta_{rt}). \quad (11)$$

Let $\theta_{r\cdot} = \{\theta_r, \theta_{\setminus r}\}$ be the set of parameters related to the node-conditional distribution of node r , where $\theta_{\setminus r} = \{\theta_{rt}\}_{t \in V \setminus r}$. Since $A_r(\cdot)$ is convex for all exponential families Wainwright and Jordan [2008], this gives us a loss function that is convex in $\theta_{r\cdot}$:

$$\ell(\theta_{r\cdot}; \mathcal{D}) = -\frac{1}{n} \sum_i \left(\left\langle B_r(X_r^{(i)}), \theta_r + \sum_{t \in V \setminus r} \theta_{rt} B_t(X_t^{(i)}) \right\rangle - A_r \left(\theta_r + \sum_{t \in V \setminus r} \theta_{rt} B_t(X_t^{(i)}) \right) \right). \quad (12)$$

We then seek to find a sparse solution in terms of both edges and individual parameters by employing the sparse group lasso regularization penalty Friedman et al. [2010], Simon et al. [2013]:

$$R(\theta_{r\cdot}) = \lambda_1 \sum_{t \in V \setminus r} \sqrt{\nu_{rt}} \|\theta_{rt}\|_2 + \lambda_2 \|\theta_{\setminus r}\|_1, \quad (13)$$

where $\nu_{rt} = m_r \times m_t$ is the number of parameters in the *pseudo-edge* from node r to node t (i.e., the edge (r, t) in the r^{th} node-conditional). This yields a collection of independent convex optimization problems, one for each node-conditional.

$$\underset{\theta_{r\cdot}}{\text{minimize}} \quad \ell(\theta_{r\cdot}; \mathcal{D}) + R(\theta_{r\cdot}) \quad (14)$$

We next present an approach to solving this problem based on Alternating Direction

Method of Multipliers (ADMM) Boyd et al. [2011].

4.1 Optimization Procedure

We first introduce a slack variable z into (14) to adhere to the canonical form of ADMM. For notational simplicity, we omit the data parameter \mathcal{D} from the loss function and the subscripts in θ_r . and A_r since it is clear we are dealing with the optimization of a single node-conditional.

$$\begin{aligned} & \underset{\theta}{\text{minimize}} && \ell(\theta) + R(z) \\ & \text{subject to} && \theta = z \end{aligned}, \quad (15)$$

where $\text{length}(\theta) = \tau$. The augmented Lagrangian is

$$L_\alpha(\theta, z, \rho) = \ell(\theta) + R(z) + \rho^T(\theta - z) + (\alpha/2) \|\theta - z\|_2^2. \quad (16)$$

Defining the residual of the slack $r = \theta - z$, we instead use the scaled form with $u = (1/\alpha)\rho$. ADMM proceeds in an alternating fashion, performing the following updates at each iteration:

$$\theta^{k+1} = \underset{\theta}{\text{argmin}} \left(\ell(\theta) + (\alpha/2) \|\theta - z^k + u^k\|_2^2 \right) \quad (17)$$

$$z^{k+1} = \underset{z}{\text{argmin}} \left(R(z) + (\alpha/2) \|\theta^{k+1} - z + u^k\|_2^2 \right) \quad (18)$$

$$u^{k+1} = u^k + \theta^{k+1} - z^{k+1} \quad (19)$$

Updating θ^{k+1} . The j^{th} subgradient of θ is $g_j(\theta) = -\bar{B}_j + \nabla_j \bar{A}(\theta) + \alpha(\theta_j + z_j^k - u_j^k)$. Note that the log-partition function, $A(\eta)$, over the natural parameters, $\eta = B\theta$, is available in closed form for most commonly-used exponential families. Thus, $\nabla^2 \bar{A}(\theta)$ is a weighted sum of rank-one matrices. In cases where the number of samples is much less than the total number of parameters (i.e. $n \ll \tau$), we can efficiently calculate an exact Newton update in $\mathcal{O}(\tau)$ by leveraging the matrix inversion lemma Boyd and Vandenberghe [2009]. Otherwise, we use a diagonal approximation of the Hessian and perform a quasi-Newton update.

Updating z^{k+1} . We can reformulate (18) as the *proximal operator* Parikh and Boyd [2013] of $R(z)$:

$$\text{prox}_{R/\alpha}(y) = \underset{z}{\text{argmin}} \left(R(z) + (\alpha/2) \|z - y\|_2^2 \right), \quad (20)$$

where $y = \theta^{k+1} + u^k$. From Friedman et al. [2010], it is straightforward to show that the update has a closed-form solution for each j^{th} block of edge parameters,

$$z_j^{k+1} = \frac{(\|S(\alpha(y_j), \lambda_2)\|_2 - \sqrt{\nu_j} \lambda_1)_+ S(\alpha(y_j), \lambda_2)}{\alpha \|S(\alpha(y_j), \lambda_2)\|_2 + \sqrt{\nu_j} \lambda_1 (1 - \alpha)}, \quad (21)$$

where $S(x, \lambda)$ is the soft-thresholding operator on x with cutoff at λ .

Updating u^{k+1} . Per ADMM, closed-form is given in (19).

We iterate each of the above update steps in turn until convergence, then AND pseudo-edges when stitching the graph back together.

4.2 Domain constraints

Many exponential family distributions require parameters with bounded domain. These bounds correspond to affine constraints on subsets of θ in the ADMM algorithm.¹ Often these constraints are simple implicit restrictions to \mathbb{R}^+ or \mathbb{R}^- . In these cases the log-normalization function $A(\eta)$ serves as a built-in *log-barrier* function. For instance, a normal distribution with unknown mean μ and unknown variance σ^2 has natural parameters $\eta_1 = \frac{\mu}{\sigma^2}$ and $\eta_2 = -\frac{1}{2\sigma^2}$, implying $\eta_2 < 0$. However, since $A(\eta) = -\frac{\eta_1^2}{4\eta_2} - \frac{1}{2} \ln(-2\eta_2)$, this constraint will be effectively enforced so long as we are given a feasible starting point for η . Such a feasible point can always be discovered using a standard phase I method Boyd and Vandenberghe [2009]. In the case of equality requirements, such as categorical and multinomial distributions, we can directly incorporate the constraints into the ADMM algorithm and solve an equality-constrained Newton’s method when updating θ .

4.3 Sparsistency

We next provide the mathematical conditions that ensure with high probability our learning procedure recovers the true graph structure underlying the joint distribution. Our results rely on similar sufficient conditions to those imposed in papers analyzing the Lasso Wainwright [2009] and the l_1/l_2 penalty in Jalali et al. [2011]. Before stating the assumptions, we introduce the notation used in the proof.

4.3.1 Notation

Let $N(r) = \{t : \theta_{rt}^* \neq 0\}$ be the true neighbourhood of node r and let d_r be the degree of r , i.e. $d_r = |N(r)|$. And S_r be the index set of parameters $\{\theta_{rj;tk}^* : t \in N(r)\}$ and similarly S_r^c be the index of parameters $\{\theta_{rj;tk}^* : t \notin N(r)\}$. From now on we will overload the notation and simply use S and S^c instead of S_r and S_r^c . Let $S_r^{(ex)} = \{\theta_{rj;tk}^* : \theta_{rj;tk}^* \neq 0 \wedge t \in N(r)\}$.

Let $Q_r^n = \nabla^2 \ell(\theta_r^*; \mathcal{D})$ be the sample Fischer Information matrix at node r . As before, we will ignore subscript r and use Q^n instead of Q_r^n . Finally, we write Q_{SS}^n for the sub-matrix indexed by S .

We use the group structured norms defined in Jalali et al. [2011] in our analysis. The group structured norm $\|u\|_{\mathcal{G}, a, b}$ of a vector u with respect to a set of disjoint groups $\mathcal{G} = \{G_1, \dots, G_T\}$ is defined as $\|(\|u_{G_1}\|_b, \dots, \|u_{G_T}\|_b)\|_a$. We ignore the group \mathcal{G} and simply use $\|u\|_{a, b}$ when it is clear from the context. Similarly the group structured

¹Note that these subsets are different than the edge-wise groups that are L_2 -penalized. Rather, these constraints apply to the sum of the i^{th} value of each edge parameter and the i^{th} bias weight.

norm $\|M\|_{(a,b),(c,d)}$ of a matrix $M_{p \times p}$ is defined as $\left\|(\|M^1\|_{c,d}, \dots, \|M^p\|_{c,d})\right\|_{a,b}$. In our analysis we always use $b = 2, d = 2$ and to minimize the notation we use $\|M\|_{a,c}$ to denote $\|M\|_{(a,2),(c,2)}$. And we define $\|M\|_{max}$ as $\max_{i,j} |M_{i,j}|$, i.e, element wise maximum of M.

4.3.2 Assumptions

Let us begin by imposing assumptions on the sample Fisher Information matrix Q^n .

Assumption 1. *Dependency condition:* $\Lambda_{min}(Q_{SS}^n) \geq C_{min}$.

Assumption 2. *Incoherence condition:*

$\|Q_{SS}^n(Q_{SS}^n)^{-1}\|_{\infty,2} \leq \frac{m_{min}}{m_{max}} \frac{(1-\alpha)}{\sqrt{d_r}}$ for some $\alpha \in (0, 1]$, where $m_{max} = \max_t m_t$, $m_{min} = \min_t m_t$.

Assumption 3. *Boundedness:*

$\Lambda_{max}(E[B(X_{V \setminus r}) B(X_{V \setminus r})^T]) \leq D_{max} < \infty$, where $B(X_{V \setminus r})$ is a vector such that $B(X_{V \setminus r}) = \{B_t(X_t)\}_{t \in V \setminus r}$.

Note that the sufficient statistics $\{B_{ri}(X_r)\}_{i=1}^{m_r}$ of node r need not be bounded. So to analyze the M-estimation problem, we make the following assumptions on log-partition functions of joint and node-conditional distributions. These are similar to the conditions imposed for sparsistency analysis of GLMs.

Assumption 4. *The log partition function of the joint distribution satisfies the following conditions: for all $r \in V$ and $i \in [m_r]$*

1. *there exists constants k_m, k_v such that $E[B_{ri}(X_r)] \leq k_m$ and $E[B_{ri}(X_r)^2] \leq k_v$,*
2. *there exists constant k_h such that $\max_{u: |u| \leq 1} \frac{\partial^2 A(\theta)}{\partial \theta_{ri}^2}(\theta_{ri}^* + u, \theta_r^*) \leq k_h$,*
3. *for scalar variable η , we define a function $\bar{A}_{r,i}$ as:*

$$\begin{aligned} \bar{A}_{r,i}(\eta; \theta) = & \log \int_{\mathcal{X}_p} \exp \left\{ \eta B_{ri}(X_r)^2 + \sum_{s \in V} C_s(X_s) \right. \\ & \left. + \sum_{s \in V} \left\langle B_s(X_s), \theta_s + \sum_{t \in N(s)} \theta_{st} B_t(X_t) \right\rangle \right\} d(x) \end{aligned} \quad (22)$$

Then, there exists a constant k_h such that $\max_{u: |u| \leq 1} \frac{\partial^2 \bar{A}_{r,i}(\eta; \theta_r^)}{\partial \eta^2}(u) \leq k_h$.*

Assumption 5. *For all $r \in V$, the log-partition function $A_r(\cdot)$ of the node wise conditional distribution satisfy that there exists functions $k_1(n, p)$ and $k_2(n, p)$ such that for all feasible pairs θ and X , $\|\nabla^2 A_r(a)\|_{max} \leq k_1(n, p)$ where $a \in [b, b + 4k_2(n, p) \max\{\log(n), \log(p)\} \mathbf{1}]$ for $b := \theta_r + \sum_{t \in V \setminus r} \theta_{rt} B_t(X_t)$, where for vectors u and v we define $[u, v] := \otimes_i [u_i, v_i]$. Moreover, we assume that $\|\nabla^3 A_r(b)\|_{max} \leq k_3(n, p)$ for all feasible pairs X and θ .*

4.3.3 Sparsistency Theorem

Given these assumptions in 4.3.2 we are now ready to state our main sparsistency result.

Theorem 2. *Consider the vector space graphical model distribution in (10) with true parameters θ^* , edge set E and vertex set V such that the assumptions 1-5 hold. Suppose that θ^* satisfies $\min_{(r,t) \in E} \|\theta_{rt}^*\|_2 \geq \frac{10 m_{max}}{C_{min}} (\lambda_1 + \lambda_2)$ and regularization parameters λ_1, λ_2 satisfy $M_1 \frac{2-\alpha}{\alpha} \frac{m_{max}}{m_{min}} \sqrt{k_1(n, p)} \sqrt{\frac{\log(pm_{max}^2)}{n}} \leq \lambda_1 + \lambda_2 \leq M_2 \frac{2-\alpha}{\alpha} k_1(n, p) k_2(n, p)$*

for positive constants M_1 and M_2 and $\lambda_2 < \left(\frac{\alpha}{2-\alpha+2 m_{max}/m_{min}} \right) \lambda_1$. Then, there exists

constants L, c_1, c_2 and c_3 such that if $n \geq \max\{L \frac{m_{max}^9}{m_{min}} d^2 k_1(n, p) (k_3(n, p))^2 (\log p')^2 \log(p m_{max}^2), \frac{4 \log(p m_{max}^2)}{k_1(n, p) k_4 k_2(n, p)^2}, \frac{8 k_h^2}{k_4^2} \log(\sum_t m_t)\}$, with probability at least $1 - c (p')^{-3} (\sum_t m_t) - \exp(-c_2 n) - \exp(-c_3 n)$, the following statements hold.

- *For each node $r \in V$, the solution of the M-estimation problem (14) is unique*
- *Moreover, for each node $r \in V$ the M-estimation problem recovers the true neighbourhood exactly.*

where $m_{max} = \max m_t$, $m_{min} = \min m_t$, $p' = \max(n, p)$.

The proof of Theorem 2 follows along similar lines to the sparsistency proof in Yang et al. [2014], albeit with a subtler analysis to support general vector-spaces. It is based on the primal dual witness proof technique and relies on the previous results. We refer the interested reader to the supplementary material for additional details regarding the proofs.

5 Experiments

We demonstrate the effectiveness of our algorithm on both synthetic data and a real-world dataset of over four million foods logged on the popular diet app, MyFitnessPal.

5.1 Synthetic experiments

The synthetic experiments were run on a vector-space mixed MRF consisting of eight Bernoulli, eight gamma (with unknown shape and rate), eight Gaussian (with unknown mean and variance), and eight Dirichlet ($k=3$) nodes. The choice of these node-conditional distributions is meant to highlight the ability of VS-MRFs to model many different types of distributions. Specifically, the Bernoulli represents a univariate, uni-parameter distribution that would still be possible to incorporate into existing mixed models. The gamma and Gaussian distributions are both multi-parameter, univariate distributions which would have required fixing one parameter (e.g. fixing the Gaussians' variances) to be compatible with previous approaches. Finally, the Dirichlet distribution is multi-parameter *and* multivariate, thereby making VS-MRFs truly unique in their ability to model this joint distribution.

For each experiment, we conducted 30 independent trials by generating random weights and sampling via Gibbs sampling with a burn-in of 2000 and thinning step size of 10. We consider two different sparsity scenarios: high (90% edge sparsity, 50% intra-edge parameter sparsity) and low (50% edge sparsity, 10% intra-edge parameter




Figure 1: ROC curves for our synthetic experiments. The top left and bottom left plots show both edge as well as within-edge-parameter recovery performance respectively, for graphs with a high degree of sparsity. The two right plots show the same performance measures, but for graphs with a relatively low degree of sparsity. The low sparsity scenario is more challenging, requiring more data to recover the majority of the graph.

sparsity). Edge recovery capability is examined by fixing λ_2 to a small value and varying λ_1 over a grid of values in the range $[0.0001, 0.5]$; parameter recovery is examined analogously by fixing λ_1 and varying λ_2 . We use AND graph stitching and measure the true positive rate (TPR) and false positive rate (FPR) as the number of samples increases from 100 to 25K.

Figure 5 shows the ROC curves at both the edge and parameter levels. The results demonstrate that our algorithm improves well as the dataset size scales. They also illustrate that graphs with a higher degree of sparsity are easier to recover with fewer samples. In both the high and low sparsity graphs, the algorithm is better able to recover the coarse-grained edge structure than the more fine-grained within-edge parameter structure, though both improve favourably with the size of the data.

5.2 MyFitnessPal Food Dataset

MyFitnessPal² (MFP) is one of the largest diet-tracking apps in the world, with over 80M users worldwide. MFP has a vast crowd-sourced database of food data, where each food entry contains a description, such as “Trader Joe’s Organic Carrots,” and a vector of sixteen macro- and micro-nutrients, such as fat and vitamin C.

²<http://myfitnesspal.com>

as:

$$P_{\text{infl}}(Z) = \exp \{ \eta_0 \mathbb{I}(Z = j) + \eta_1^T B(z) + C(z) - A_{\text{infl}}(\eta) \},$$

where $A_{\text{infl}}(\eta)$ is the log-normalization constant of the point-inflated distribution which can be expressed in terms of the log-normalization constant $A(\cdot)$ of the uninflated exponential family distribution: $A_{\text{infl}}(\eta) = \log (\exp\{\eta_0\} - \exp\{\eta_1^T B(j)\mathbb{I}(j \in \mathcal{Z})\} + \exp\{A(\eta_1)\})$. Thus, as long as we have a closed form $A(\cdot)$ for the log-partition function of the base distribution, we can efficiently calculate $P_{\text{infl}}(Z)$. The definition also permits an arbitrary number of inflated points by recursively specifying the base distribution as another point-inflated model. We model each of the MFP nutrients via a two-point-inflated gamma distribution, with points at zero and a winsorized outlier bucket.

Due to the size of the overall graph, presenting it in closer detail here is not feasible. To give a qualitative perspective of the relationships captured by our algorithm, we selected the top 100 edges in the MFP food graph by ranking the edges based on their L2-norm. We then calculated their approximate contribution to the mean of their corresponding node-conditionals to determine edge color and thickness. Figure 2 shows the results of this process, with edges that contribute positively colored in orange and edges that contribute negatively colored in blue; edge thickness corresponds to the magnitude of the contribution. A high-level view of the entire learned graph is available in the supplementary materials.

Several interesting relationships can be discovered, even from just this small subset of the overall model. For instance, the negative connection between “peanut” and sodium may seem counter-intuitive, given the popularity of salted nuts, yet on inspection of the raw database it appears that indeed many peanut-based foods are actually very low in sodium on a per-calorie basis. As another example, “chips” are often thought of as a high-carb food, but the graph suggests that they actually tend to be a bigger indicator of high fat. In general, we believe there is great potential for wide-ranging future uses of VS-MRFs in nutrition and other scientific fields, with the MFP case study only scratching the surface of what can be achieved.

6 Conclusion

We have presented vector-space MRFs as a flexible and scalable approach to modeling complex, heterogeneous data. In particular, we generalize the concept of mixed MRFs to allow for node-conditional distributions to be distributed according to a generic exponential family distribution, that is potentially multi-parameter and even multivariate. Our VS-MRF learning algorithm has reassuring sparsistency guarantees and was validated against a variety of synthetic experiments and a real-world case study. We believe that the broad applicability of VS-MRFs will make them a valuable addition to the scientific toolbox. All code for our VS-MRF implementation is publicly available.⁵

⁵<https://github.com/tansey/vsmrfs>

A Proof of Theorem 1

The proof follows the same lines as the proof in Yang et al. [2014]. Let us denote $Q(X)$ as $\log(P(X)/P(0))$. Note that $X = (X_1, X_2, \dots, X_p)$ and each X_r belongs to a vector space. Given any X , let us denote \bar{X}_s as $\bar{X}_s = (X_1, \dots, X_{s-1}, 0, X_{s+1}, \dots, X_p)$. Consider the following expansion for $Q(X)$:

$$\begin{aligned} Q(X) = & \sum_{t \in \{1, \dots, p\}} \mathcal{I}[X_t \neq 0] G_t(X_t) + \dots \\ & + \sum_{\substack{t_1, \dots, t_k \in \\ \{1, \dots, p\}}} \mathcal{I}[X_{t_1} \neq 0, \dots, X_{t_k} \neq 0] G_{t_1 \dots t_k}(X_{t_1} \dots X_{t_k}) \end{aligned} \quad (23)$$

where \mathcal{I} is the indicator function which takes value 1 if its argument evaluates to true and 0 otherwise.

Using some simple algebra and the definition $Q(X) = \log(P(X)/P(0))$ we can show that

$$\exp(Q(X) - Q(\bar{X}_s)) = \frac{P(X_s | X_1, \dots, X_{s-1}, X_{s+1}, \dots, X_p)}{P(0 | X_1, \dots, X_{s-1}, X_{s+1}, \dots, X_p)} \quad (24)$$

From (23) we have the following:

$$\begin{aligned} (Q(X) - Q(\bar{X}_s)) = & \mathcal{I}[X_s \neq 0] \left(G_s(X_s) + \sum_{t \in \{1, \dots, p\} \setminus s} \mathcal{I}[X_t \neq 0] G_{s,t}(X_s, X_t) \right. \\ & \left. + \sum_{\substack{t_2, \dots, t_k \in \\ \{1, \dots, p\} \setminus s}} \mathcal{I}[X_{t_2} \neq 0, \dots, X_{t_k} \neq 0] G_{s, t_2 \dots t_k}(X_s, \dots, X_{t_k}) \right) \end{aligned} \quad (25)$$

Since the node conditional distribution follows the exponential family distribution defined in (6) we can show that:

$$\log \frac{P(X_s | X_1, \dots, X_{s-1}, X_{s+1}, \dots, X_p)}{P(0 | X_1, \dots, X_{s-1}, X_{s+1}, \dots, X_p)} = \langle E_s(X_{-s}), B_s(X_s) - B_s(0) \rangle + (C_s(X_s) - C_s(0)) \quad (26)$$

Using (25) and (26) for left and right hand sides of (24) and setting $X_t = 0$ for all $t \neq s$ we obtain:

$$\mathcal{I}[X_s \neq 0] G_s(X_s) = \langle E_s(0), B_s(X_s) - B_s(0) \rangle + (C_s(X_s) - C_s(0))$$

Similarly setting $X_r = 0$ for all $r \notin \{s, t\}$ we obtain:

$$\begin{aligned} \mathcal{I}[X_s \neq 0]G_s(X_s) + \mathcal{I}[X_s \neq 0, X_t \neq 0]G_{s,t}(X_s, X_t) = \\ \langle E_s(0 \cdots X_t \cdots, 0), B_s(X_s) - B_s(0) \rangle + (C_s(X_s) - C_s(0)) \end{aligned}$$

Similarly, replacing X_s with X_t in (24) and setting $X_r = 0$ for all $r \notin \{s, t\}$ we obtain:

$$\begin{aligned} \mathcal{I}[X_t \neq 0]G_t(X_t) + \mathcal{I}[X_s \neq 0, X_t \neq 0]G_{s,t}(X_s, X_t) = \\ \langle E_t(0 \cdots X_s \cdots, 0), B_t(X_t) - B_t(0) \rangle + (C_t(X_t) - C_t(0)) \end{aligned}$$

From the above three equations we arrive at the following equality:

$$\begin{aligned} \langle E_s(0 \cdots X_t \cdots, 0) - E_s(0), B_s(X_s) - B_s(0) \rangle = \\ \langle E_t(0 \cdots X_s \cdots, 0) - E_t(0), B_t(X_t) - B_t(0) \rangle \end{aligned} \quad (27)$$

The above equality should hold for the node conditional distributions to be consistent with the joint MRF distribution over X with respect to graph G . So we need to find the form of $E_r(\cdot)$ that satisfies the above equation. Omitting zero vectors for clarity from (27), we get the following:

$$\begin{aligned} \langle E_t(X_s), B_t(X_t) \rangle &= \langle E_s(X_t), B_s(X_s) \rangle \\ \sum_j E_{tj}(X_s) B_{tj}(X_t) &= \sum_l E_{sl}(X_t) B_{sl}(X_s) \end{aligned} \quad (28)$$

We rewrite the natural parameter functions as

$$\begin{aligned} E_{tj}(X_s) &= \sum_l \theta_{sl;tj} B_{sl}(X_s) + \bar{B}_{tj}(X_s) \\ E_{sl}(X_t) &= \sum_j \bar{\theta}_{sl;tj} B_{tj}(X_t) + \bar{B}_{sl}(X_t) \end{aligned} \quad (29)$$

where $\forall j \bar{B}_{tj}(X_s)$ are functions in the Hilbert space \mathcal{H}_s orthogonal to the span of functions $B_s(X_s)$, and $\forall j \bar{B}_{sl}(X_t)$ are functions in the Hilbert space \mathcal{H}_t orthogonal to the span of functions $B_t(X_t)$; and $\theta_{sl;tj}, \bar{\theta}_{sl;tj}$ are scalars. Combining (28) and (29), we get

$$\begin{aligned} \sum_j \sum_l \theta_{sl;tj} B_{sl}(X_s) B_{tj}(X_t) + \sum_j \bar{B}_{tj}(X_s) B_{tj}(X_t) \\ = \sum_l \sum_j \bar{\theta}_{sl;tj} B_{sl}(X_s) B_{tj}(X_t) + \sum_l \bar{B}_{sl}(X_t) B_{sl}(X_s) \end{aligned} \quad (30)$$

Rearranging terms in the above equation gives us the following equation:

$$\begin{aligned} & \sum_j \left(\sum_l (\theta_{sl;tj} - \bar{\theta}_{sl;tj}) B_{sl}(X_s) + \bar{B}_{tj}(X_s) \right) B_{tj}(X_t) \\ &= \sum_l B_{sl}(X_s) \bar{B}_{sl}(X_t) \end{aligned} \quad (31)$$

However, since $\forall l \bar{B}_{sl}(X_t)$ is orthogonal to $B_t(X_t)$, the left and right hand sides of the above equation are equal to 0, which leads us to the following equations.

$$\begin{aligned} & \sum_l B_{sl}(X_s) \bar{B}_{sl}(X_t) = 0 \\ & \sum_j \left(\sum_l (\theta_{sl;tj} - \bar{\theta}_{sl;tj}) B_{sl}(X_s) + \bar{B}_{tj}(X_s) \right) B_{tj}(X_t) = 0 \end{aligned} \quad (32)$$

However since we assumed that the sufficient statistics are minimal we get $\forall l \bar{B}_{sl}(X_t) = 0$ from the first equality and $\forall j, l \theta_{sl;tj} = \bar{\theta}_{sl;tj}, \bar{B}_{tj}(X_s) = 0$ from the second equality.

Hence from (29), we obtain $E_s(X_t) = \theta_{st}(B_t(X_t) - B_t(0))$ and $E_t(X_s) = \theta_{st}^T(B_s(X_s) - B_s(0))$ where θ_{st} is a matrix formed by the scalars $\theta_{sl;tj}$ such that $(\theta_{st})_{lj} = \theta_{sl;tj}$ and:

$$\begin{aligned} & \mathcal{I}[X_s \neq 0, X_t \neq 0] G_{s,t}(X_s, X_t) = \\ & (B_t(X_t) - B_t(0))^T \theta_{st}^T (B_s(X_s) - B_s(0)) \end{aligned} \quad (33)$$

By extending this argument to higher order factors we can show that the natural parameters are required to be in the form specified by (7).

B Proof of Sparsistency

Before proving the sparsistency result, we will show that the sufficient statistics $B_r(X_r)$ are well behaved. Recall that $B_{ri}(X_r)$ indicates i^{th} component of the vector $B_r(X_r)$. We set the convention that whenever a variable has the subscript $\setminus r$ attached we will be referring to the set of indexes $\{(t, j, k) : \theta_{rj;tk} \in \theta_{r\cdot}, t \neq r\}$.

Proposition 1. *Let $\{X^{(j)}\}_{j=1}^n$ have joint distribution as in (10), then,*

$$P \left(\frac{1}{n} \sum_{j=1}^n \left(B_{ri} \left(X_r^{(j)} \right) \right)^2 \geq \delta \right) \leq \exp \left(-n \frac{\delta^2}{4k_h^2} \right) \quad (34)$$

for $\delta \leq \min \{2 \frac{k_v}{3}, k_h + k_v\}$.

Proof. It is clear from Taylor Series expansion and assumption 4 that

$$\begin{aligned}
& \log E \left[\exp \left(t B_{ri} (X_r)^2 \right) \right] = \\
& \log \int_{\otimes_{s \in [p]} \mathcal{X}_s} \exp \left\{ t B_{ri} (X_r)^2 + \right. \\
& \sum_{s \in V} \left\langle B_s (X_s), \theta_s^* + \sum_{t \in N(r)} \theta_{st}^* B_t (X_t) \right\rangle + \\
& \left. \sum_{s \in V} C_s (X_s) - A(\theta^*) \right\} v(dx) \\
& = \bar{A}_{r,i}(\eta; \theta)(t; \theta^*) - \bar{A}_{r,i}(\eta; \theta)(0; \theta^*) \\
& \leq t \frac{\partial \bar{A}_{r,i}(\eta; \theta)}{\partial \eta}(0) + \frac{t^2}{2} \frac{\partial^2 \bar{A}_{r,i}(\eta; \theta)}{\partial \eta^2}(ut) \\
& \leq t k_v + \frac{t^2}{2} k_h
\end{aligned} \tag{35}$$

where $u \in [0, 1]$

Therefore, by the standard Chernoff bounding technique, for $t \leq 1$, it follows that

$$\begin{aligned}
& \mathbb{P} \left(\frac{1}{n} \sum_{j=1}^n B_{ri} \left(X_r^{(j)} \right)^2 \geq \delta \right) \leq \\
& \exp \left(-n \delta t + n k_v t + \frac{t^2}{2} K_h n \right) \leq \\
& \exp \left(-n \frac{\delta^2}{4 k_h^2} \right)
\end{aligned} \tag{36}$$

for $\delta \leq \min \left\{ 2 \frac{k_v}{3}, k_h + k_v \right\}$. \square

Proposition 2. Let X be a random vector with the distribution specified in (10). Then, for any positive constant δ and some constant $c > 0$

$$P \left(|B_{ri}(X_r)| \geq \delta \log(\eta) \right) \leq c \eta^{-\delta} \tag{37}$$

Proof. Let \bar{v} be a unit vector with the same dimensions as θ_r^* , and exactly one non-zero entry, corresponding to the sufficient statistic $B_{ri}(X_r)$. Then we can write $\log \left(E[\exp(B_{ri}(X_r))] \right)$ as:

$$\log \left(E[\exp(B_{ri}(X_r))] \right) = A(\theta^* + \bar{v}) - A(\theta^*)$$

By Taylor series expansion, for some $u \in [0, 1]$, we can rewrite last equation as

$$\begin{aligned}
A(\theta^* + \bar{v}) - A(\theta^*) &= \nabla A(\theta^*) \cdot \bar{v} + \frac{1}{2} \bar{v}^T \nabla^2 A(\theta^* + u \bar{v}) \bar{v} \\
&= E[B_{ri}(X_r)] \|\bar{v}\|_2 \\
&\quad + \frac{1}{2} \frac{\partial^2 A(\theta^* + u \bar{v})}{\partial \theta_{ri}^2} \|\bar{v}\|_2^2
\end{aligned}$$

Using Assumption 4 we get the inequality :

$$A(\theta^* + \bar{v}) - A(\theta^*) \leq k_m + \frac{1}{2}k_h$$

Now, by using Chernoff bound, for any positive constant a , we get $P(B_{ri}(X_r) \geq a) \leq \exp(-a + k_m + \frac{1}{2}k_h)$. By setting $a = \delta \log(\eta)$ it follows that

$$P(B_{ri}(X_r) \geq \delta \log(\eta)) \leq \exp(-\delta \log(\eta) + k_m + \frac{1}{2}k_h) \leq c\eta^{-\delta}$$

where $c = \exp(k_m + \frac{1}{2}k_h)$ □

The proof of Sparsistency is based on the primal dual witness proof technique. First note that the optimality condition of (14), can be written as:

$$\nabla \ell(\hat{\theta}_r; \mathcal{D}) + \lambda_1 \sum_{t:r \neq t} \sqrt{\nu_{rt}} \hat{Z}_{1,rt} + \lambda_2 \hat{Z}_2 = 0 \quad (38)$$

where $\hat{Z}_{1,rt} \in \partial \|\hat{\theta}_{rt}\|_2$, $\hat{Z}_2 \in \partial \|\hat{\theta}_{\setminus r}\|_1$ and we denote $\hat{Z} = (\hat{Z}_1, \hat{Z}_2)$, where $\hat{Z}_1 = \{\hat{Z}_{1,rt}\}_{t \in V \setminus r}$. And sub-gradients \hat{Z}_1, \hat{Z}_2 should satisfy the following conditions:

$$\begin{aligned} \forall i \quad (\hat{Z}_2)_i &= \text{sign}((\hat{\theta}_r)_i) \quad \text{if } (\hat{\theta}_r)_i \neq 0 \\ &|(\hat{Z}_2)_i| \leq 1 \quad \text{otherwise} \\ \forall t \quad \hat{Z}_{1,rt} &= \frac{\hat{\theta}_{rt}}{\|\hat{\theta}_{rt}\|_2} \quad \text{if } \hat{\theta}_{rt} \neq 0 \\ &\|\hat{Z}_{1,rt}\|_2 \leq 1 \quad \text{otherwise} \end{aligned} \quad (39)$$

Note that we can think of \hat{Z}_1 and \hat{Z}_2 as dual variables by appealing to Lagrangian theory. The next lemma shows that graph structure recovery is guaranteed if the dual is strictly feasible.

Lemma 1. *Suppose that there exists a primal-dual pair $(\hat{\theta}_r, \hat{Z})$ for (14) such that $\|\hat{Z}_{1,S^c}\|_{\infty,2} < 1$ and $\|\hat{Z}_{2,S^c}\|_{\infty} < 1$. Then, any optimal solution $\tilde{\theta}_r$ must satisfy $(\tilde{\theta}_r)_{S^c} = 0$. Moreover, if the Hessian sub-matrix $[\nabla^2 \ell(\hat{\theta}_r)]_{SS}$ is positive definite then $\hat{\theta}_{\setminus r}$ is the unique optimal solution.*

Proof. First, note that by Cauchy–Schwarz’s and Holder’s inequalities

$$\langle \hat{Z}_{1,rt}, \tilde{\theta}_{rt} \rangle \leq \|\tilde{\theta}_{rt}\|_2 \quad \text{and} \quad \langle \hat{Z}_2, \tilde{\theta}_{\setminus r} \rangle \leq \|\tilde{\theta}_{\setminus r}\|_1. \quad (40)$$

But from (38) and the primal optimality of $\hat{\theta}_{r\cdot}$ and $\tilde{\theta}_{r\cdot}$ for (14),

$$\begin{aligned}
& \ell(\tilde{\theta}_{r\cdot}) + \sum_{t \neq r} \lambda_1 \sqrt{\nu_{rt}} \langle \hat{Z}_{1,rt}, \tilde{\theta}_{rt} \rangle + \lambda_2 \langle \hat{Z}_2, \tilde{\theta}_{\setminus r} \rangle \\
& \geq \min_{\theta} \ell(\theta_{r\cdot}) + \sum_{t \neq r} \lambda_1 \sqrt{\nu_{rt}} \langle \hat{Z}_{1,rt}, \theta_{rt} \rangle + \lambda_2 \langle \hat{Z}_2, \theta_{\setminus r} \rangle \\
& = \ell(\hat{\theta}_{r\cdot}) + \sum_{t \neq r} \lambda_1 \sqrt{\nu_{rt}} \langle \hat{Z}_{1,rt}, \hat{\theta}_{rt} \rangle + \lambda_2 \langle \hat{Z}_2, \hat{\theta}_{\setminus r} \rangle \\
& = \ell(\tilde{\theta}_{r\cdot}) + \sum_{t \neq r} \lambda_1 \sqrt{\nu_{rt}} \|\tilde{\theta}_{rt}\|_2 + \|\tilde{\theta}_{\setminus r}\|_1
\end{aligned} \tag{41}$$

hence, combining with (40) with (41) it follows that $\sum_{t \neq r} \lambda_1 \sqrt{\nu_{rt}} \|\tilde{\theta}_{rt}\|_2 + \|\tilde{\theta}_{\setminus r}\|_1 = \sum_{t \neq r} \lambda_1 \sqrt{\nu_{rt}} \|\hat{\theta}_{rt}\|_2 + \|\hat{\theta}_{\setminus r}\|_1 = \sum_{t \neq r} \lambda_1 \sqrt{\nu_{rt}} \langle \hat{Z}_{1,rt}, \hat{\theta}_{rt} \rangle + \lambda_2 \langle \hat{Z}_2, \hat{\theta}_{\setminus r} \rangle$. The result follows.

If the Hessian sub-matrix is positive definite for the restricted problem then the problem is strictly convex and has a unique solution. \square

Based on the above lemma, we prove sparsistency theorem by constructing a primal-dual witness $(\hat{\theta}_{r\cdot}, \hat{Z})$ with the following steps:

1. Set $(\hat{\theta}_{r\cdot})_S = \operatorname{argmin}_{((\theta_{r\cdot})_S, 0)} \ell((\theta_{r\cdot})_S; \mathcal{D}) + \lambda_1 \sum_{t \in S} \sqrt{\nu_{rt}} \|\theta_{rt}\|_2 + \lambda_2 \|(\theta_{r\cdot})_S\|_1$
2. For $t \in S$, we define $\hat{Z}_{1,rt} = \frac{\theta_{rt}}{\|\theta_{rt}\|_2}$ and then construct $\hat{Z}_{2,S}$ by the stationary condition.
3. Set $(\hat{\theta}_{r\cdot})_{S^c} = 0$
4. Set \hat{Z}_{2,S^c} such that $\|\hat{Z}_{2,S^c}\|_\infty < 1$
5. Set \hat{Z}_{1,S^c} such that condition (38) is satisfied.
6. The final step consists of showing, that the following conditions are satisfied:
 - (a) *strict dual feasibility* : the condition in Lemma 1 holds with high probability
 - (b) *correct neighbourhood recovery*: the primal-dual pair specifies the neighbourhood of r , with high probability

We begin by proving some key lemmas that are key to our main theorem. The sub-gradient optimality condition (38) can be rewritten as:

$$\nabla \ell(\hat{\theta}_{r\cdot}; \mathcal{D}) - \nabla \ell(\theta_{r\cdot}^*; \mathcal{D}) = W^n - \lambda_1 \sum_{t:r \neq t} \sqrt{\nu_{rt}} \hat{Z}_{1,rt} - \lambda_2 \hat{Z}_2 \tag{42}$$

where $W^n = -\nabla \ell(\theta_{r\cdot}^*; \mathcal{D})$ and $\theta_{r\cdot}^*$ is the true model parameter. By applying mean-value theorem coordinate wise to (42), we get:

$$\begin{aligned}
\nabla^2 \ell(\theta_{r\cdot}^*; \mathcal{D}) [\hat{\theta}_{r\cdot} - \theta_{r\cdot}^*] &= W^n - \lambda_1 \sum_{t:r \neq t} \sqrt{\nu_{rt}} \hat{Z}_{1,rt} \\
&\quad - \lambda_2 \hat{Z}_2 + R^n
\end{aligned} \tag{43}$$

where R^n is the remainder term after applying mean-value theorem: $R_j^n = [\nabla^2 \ell(\theta_r^*; \mathcal{D}) - \nabla^2 \ell(\bar{\theta}_{r.}^j; \mathcal{D})]_j^T (\hat{\theta}_{r.} - \theta_r^*)$ for some $\bar{\theta}_{r.}^j$ on the line between $\hat{\theta}_{r.}$ and θ_r^* , and with $[\cdot]_j^T$ denoting the j -th row of matrix. The following lemma controls the score term W^n

Lemma 2. Recall $\nu_{r \max} = \max_t \nu_{rt}$, $\nu_{r \min} = \min_t \nu_{rt}$, $p' = \max(n, p)$. Assume that

$$\begin{aligned} \frac{8(2-\alpha)}{\alpha} \sqrt{k_1(n, p) k_4 \frac{\nu_{r \max} \log(p \nu_{r \max})}{n \nu_{r \min}}} &\leq \\ \lambda_1 + \lambda_2 &\leq \frac{4(2-\alpha) \sqrt{\nu_{r \max}}}{\alpha \sqrt{\nu_{r \min}}} k_1(n, p) k_2(n, p) k_4 \end{aligned} \quad (44)$$

for some constant $k_4 \leq \min \{2 \frac{k_v}{3}, k_h + k_v\}$ and suppose also that $n \geq \frac{8 k_h^2}{k_4^2} \log(\sum_t m_t)$ then,

$$\begin{aligned} P \left(\|W_{\setminus r}^n\|_{\infty, 2} > \frac{\alpha}{2-\alpha} \frac{\sqrt{\nu_{r \min}} (\lambda_1 + \lambda_2)}{4} \right) &\leq \\ 1 - c_1 p'^{-3} (\sum_t m_t) - \exp(-c_2 n) - \exp(-c_3 n) \end{aligned} \quad (45)$$

Proof. Define $W_t^n = -\nabla_{\theta_{rt}} \ell(\theta_r^*; \mathcal{D})$. Let $W_{t,jk}^n$ be the element in W_t^n corresponding to parameter $\theta_{rj;tk}$. Note that $W_{t,jk}^n = \frac{1}{n} \sum_{i=1}^n V_{t,jk}^i$ where

$$\begin{aligned} V_{t,jk}^i &= B_{rj} \left(X_r^{(i)} \right) B_{tk} \left(X_t^{(i)} \right) - \\ &\quad \nabla_{\theta_{rj;tk}} A_r \left(\theta_r^* + \sum_{s \in V \setminus r} \theta_{rs}^* B_s \left(X_s^{(i)} \right) \right) B_{tk} \left(X_t^{(i)} \right) \end{aligned}$$

so for $t' \in \mathbb{R}$

$$\begin{aligned} E \left[\exp \left(t' V_{t,jk}^i \right) | X_{V \setminus r}^{(i)} \right] &= \\ \int_{X_r^{(i)}} \exp \left\{ t' \left[B_{rj} \left(X_r^{(i)} \right) B_{tk} \left(X_t^{(i)} \right) \right. \right. \\ &\quad \left. \left. - \nabla_{\theta_{rj;tk}} A_r \left(\theta_r^* + \sum_{s \in V \setminus r} \theta_{rs}^* B_s \left(X_s^{(i)} \right) \right) B_{tk} \left(X_t^{(i)} \right) \right] \right. \\ &\quad \left. + C \left(X_r^{(i)} \right) + \theta_r^* B_r \left(X_r^{(i)} \right) + \sum_{s \in V \setminus r} B_r \left(X_r^{(i)} \right) \theta_{rs}^* B_s \left(X_s^{(i)} \right) \right. \\ &\quad \left. - A_r \left(\theta_r^* + \sum_{s \in V \setminus r} \theta_{rs}^* B_s \left(X_s^{(i)} \right) \right) \right\} dX_r \\ &= \exp \left\{ A_r \left(\theta_r^* + t' B_{tk} \left(X_t^{(i)} \right) + \sum_{s \in V \setminus r} \theta_{rs}^* B_s \left(X_s^{(i)} \right) \right) \right. \\ &\quad \left. - A_r \left(\theta_r^* + \sum_{s \in V \setminus r} \theta_{rs}^* B_s \left(X_s^{(i)} \right) \right) \right. \\ &\quad \left. - \nabla_{\theta_{rj;tk}} A_r \left(\theta_r^* + \sum_{s \in V \setminus r} \theta_{rs}^* B_s \left(X_s^{(i)} \right) \right) t' B_{tk} \left(X_t^{(i)} \right) \right\} \\ &= \exp \left\{ \frac{\nabla_{\theta_{rj;tk}}^2 A_r(c)}{2} B_{tk} \left(X_t^{(i)} \right)^2 t'^2 \right\} \end{aligned}$$

where $c = \theta_r^* + \sum_{s \neq r} \theta_{rs}^* B_s \left(X_s^{(i)} \right) + v_1 t' B_{tk} \left(X_t^{(i)} \right)$ for some $v_1 \in [0, 1]$. Therefore,

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \log E \left[\exp(t' V_{t,jk}^i) | X_{V \setminus r}^{(i)} \right] = \\ & \frac{1}{n} \sum_{i=1}^n \frac{1}{2} \nabla_{\theta_{rj;tk}, \theta_{rj;tk}}^2 A_r(c) B_{tk} \left(X_t^{(i)} \right)^2 t'^2 \end{aligned}$$

Next lets define event $\varepsilon_1 = \{ \max_{i,t} \|B_t(X_t^{(i)})\|_\infty \leq 4 \log p' \}$. Then, from Proposition 2 we get $P(\varepsilon_1^c) \leq c_1 n p'^{-4} (\sum_t m_t) \leq c_1 p'^{-3} (\sum_t m_t)$. If $t' \leq k_2(n, p)$, Assumption 5 implies that

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \log E \left[\exp(t' V_{t,jk}^i) | X_{V \setminus r}^{(i)} \right] \leq \\ & \frac{k_1(n, p)}{2} \frac{1}{n} \sum_{i=1}^n B_{tk} \left(X_t^{(i)} \right)^2 t'^2 \end{aligned}$$

Now, lets define event $\varepsilon_2 = \{ \max_{t,j} \frac{1}{n} \sum_{i=1}^n \left(B_{tj} \left(X_t^{(i)} \right) \right)^2 \leq k_4 \}$ where $k_4 \leq \min\{2k_v/3, k_h + k_v\}$. Then, by proposition (1) we obtain that if $n \geq \frac{8k_h^2}{k_4^2} \log(\sum_{t \in V} m_t)$:

$$P(\varepsilon_2^c) \leq \exp \left(-n \frac{k_4^2}{4k_h^2} + \log \left(\sum_{t \in V} m_t \right) \right) \leq \exp(-n c_2) \quad (46)$$

Therefore, for $t' \leq k_2(n, p)$,

$$\frac{1}{n} \sum_{i=1}^n \log E \left[\exp(t' V_{t,jk}^i) | X_{V \setminus r}^{(i)} \right] \leq \frac{k_1(n, p) k_4 t'^2}{2} \quad (47)$$

Hence, by the standard Chernoff bound technique, for $t' \leq k_2(n, p)$

$$\begin{aligned} & P \left(\frac{1}{n} \sum_{i=1}^n |V_{t,jk}^i| > \delta \mid \varepsilon_1, \varepsilon_2 \right) \leq \\ & 2 \exp \left(n \left(\frac{k_1(n, p) k_4 t'^2}{2} - t' \delta \right) \right) \end{aligned} \quad (48)$$

Setting $t' = \frac{\delta}{k_1(n, p) k_4}$, for $\delta \leq k_1(n, p) k_2(n, p) k_4$, we arrive to:

$$P \left(\frac{1}{n} \sum_{i=1}^n |V_{t,jk}^i| > \delta \mid \varepsilon_1, \varepsilon_2 \right) \leq 2 \exp \left(\frac{-n \delta^2}{2 k_1(n, p) k_4} \right) \quad (49)$$

Supposing that $\frac{\alpha \sqrt{\nu_r \min}}{2-\alpha} \frac{\lambda_1 + \lambda_2}{4 \sqrt{m_r m_t}} \leq k_1(n, p) k_2(n, p) k_4$. It then follows that $\delta = \frac{\alpha \sqrt{\nu_r \min}}{2-\alpha} \frac{\lambda_1 + \lambda_2}{4 \sqrt{m_r m_t}}$ satisfies

$$P \left(\frac{1}{n} \sum_{i=1}^n |V_{t,jk}^i| > \frac{\alpha \sqrt{\nu_{r \min}}}{2-\alpha} \frac{\lambda_1 + \lambda_2}{4 \sqrt{m_r m_t}} \mid \varepsilon_1, \varepsilon_2 \right) \leq 2 \exp \left(\frac{-\alpha^2}{(2-\alpha)^2} \frac{\nu_{r \min} n (\lambda_1 + \lambda_2)^2}{32 k_1(n,p) k_4 m_r m_t} \right) \quad (50)$$

Form which, we obtain the following using union bound

$$\begin{aligned} P \left(\|W_t^n\|_2 > \frac{\alpha \sqrt{\nu_{r \min}} (\lambda_1 + \lambda_2)}{2-\alpha} \mid \varepsilon_1, \varepsilon_2 \right) &\leq \\ P \left(\|W_t^n\|_\infty > \frac{\alpha \sqrt{\nu_{r \min}} (\lambda_1 + \lambda_2)}{2-\alpha} \mid \varepsilon_1, \varepsilon_2 \right) & \\ \leq 2 \exp \left(\frac{-\alpha^2}{(2-\alpha)^2} \frac{\nu_{r \min} n (\lambda_1 + \lambda_2)^2}{32 k_1(n,p) k_4 m_r m_t} + \log(\nu_{rt}) \right) & \end{aligned} \quad (51)$$

and hence,

$$\begin{aligned} P \left(\|W^n\|_{\infty,2} > \frac{\alpha \sqrt{\nu_{r \min}} (\lambda_1 + \lambda_2)}{2-\alpha} \mid \varepsilon_1, \varepsilon_2 \right) &\leq \\ 2 \exp \left(\frac{-\alpha^2}{(2-\alpha)^2} \frac{\nu_{r \min} n (\lambda_1 + \lambda_2)^2}{32 k_1(n,p) k_4 \nu_{r \max}} + \log(\nu_{r \max}) + \log p \right) & \end{aligned} \quad (52)$$

Finally for $\lambda_1 + \lambda_2 \geq \frac{8(2-\alpha)}{\alpha} \sqrt{k_1(n,p) k_4 \frac{\nu_{r \max} \log(p \nu_{r \max})}{n \nu_{r \min}}}$, we obtain

$$\begin{aligned} P \left(\|W^n\|_{\infty,2} > \frac{\alpha \sqrt{\nu_{r \min}} (\lambda_1 + \lambda_2)}{2-\alpha} \right) &\leq \\ c_1 p'^{-3} (\sum_t m_t) + \exp(-c_2 n) + \exp(-c_3 n) & \end{aligned} \quad (53)$$

□

Lemma 3. Suppose that $\lambda_1 + \lambda_2 \leq \frac{C_{\min}^2}{40 \log p' D_{\max} d_r k_3(n,p) \nu_{r \max}^2}$ and $\|W_{\setminus r}^n\|_{\infty,2} \leq \frac{(\lambda_1 + \lambda_2) \alpha \sqrt{\nu_{r \min}}}{4(2-\alpha)}$, then,

$$\begin{aligned} P \left(\|(\theta_{r \cdot}^*)_S - (\hat{\theta}_{r \cdot})_S\|_{\infty,2} \leq \frac{5 \sqrt{\nu_{r \max}}}{C_{\min}} (\lambda_1 + \lambda_2) \right) & \\ \geq 1 - c p'^{-3} (\sum_t m_t) & \end{aligned} \quad (54)$$

for some constant $c > 0$.

Proof. We define $F(u_S)$ as:

$$\begin{aligned} F(u_S) = & \ell((\theta_{r \cdot}^*)_S + u_S; \mathcal{D}) - \ell((\theta_{r \cdot}^*)_S; \mathcal{D}) \\ & + \lambda_1 \sum_{t \in N(r)} \sqrt{\nu_{rt}} (\|\theta_{rt}^* + u_{rt}\|_2 - \|\theta_{rt}^*\|_2) \\ & + \lambda_2 (\|(\theta_{r \cdot}^*)_S + u_S\|_1 - \|(\theta_{r \cdot}^*)_S\|_1) \end{aligned} \quad (55)$$

From the construction of $\hat{\theta}_{r \cdot}$, it is clear that $\hat{u}_s = (\hat{\theta}_{r \cdot})_S - (\theta_{r \cdot}^*)_S$ minimizes F . And since $F(0) = 0$, we have $F(\hat{u}_s) \leq 0$. We now show that for some $B > 0$ with $\|u_S\|_{\infty,2} = B$, we have $F(u_S) > 0$. Using this and the fact that F is convex we can then show that $\|\hat{u}_s\|_{\infty,2} \leq B$.

Let u_S an arbitrary vector with $\|u_S\|_{\infty,2} = \frac{5 \sqrt{\nu_{r \max}}}{C_{\min}} (\lambda_1 + \lambda_2)$. Then, from the Taylor Series expansion of log likelihood function in F , we have:

$$\begin{aligned}
F(u_S) = & \nabla \ell((\theta_{r\cdot}^*)_S; \mathcal{D})^T u_S \\
& + u_S^T \nabla^2 \ell((\theta_{r\cdot}^*)_S + v u_S) u_S \\
& + \lambda_1 \sum_{t \in N(r)} \sqrt{\nu_{rt}} (\|\theta_{rt}^* + u_{rt}\|_2 - \|\theta_{rt}^*\|_2) \\
& + \lambda_2 (\|(\theta_{r\cdot}^*)_S + u_S\|_1 - \|(\theta_{r\cdot}^*)_S\|_1)
\end{aligned} \tag{56}$$

for some $v \in [0, 1]$

We now bound each of the terms in the right hand side of (56). From (51) and using Cauchy-Schwarz inequality we obtain:

$$\begin{aligned}
& \left| \nabla \ell((\theta_{r\cdot}^*)_S; \mathcal{D})^T u_S \right| \\
& \leq \|\nabla \ell((\theta_{r\cdot}^*)_S; \mathcal{D})\|_\infty \|u_S\|_1 \\
& \leq \|\nabla \ell((\theta_{r\cdot}^*)_S; \mathcal{D})\|_\infty d_r \sqrt{\nu_{r \max}} \|u_S\|_{\infty, 2} \\
& \leq \frac{\alpha}{2-\alpha} \frac{\lambda_1 + \lambda_2}{4} d_r \nu_{r \max} \frac{5}{C_{\min}} (\lambda_1 + \lambda_2) \\
& = \frac{5 \nu_{r \max}}{4 C_{\min}} d_r (\lambda_1 + \lambda_2)^2
\end{aligned} \tag{57}$$

where the last inequality holds because $\alpha \in (0, 1]$. Moreover, from triangle inequality we have:

$$\begin{aligned}
& \lambda_1 \sum_{t \in S} \sqrt{\nu_{rt}} (\|\theta_{rt}^* + u_{rt}\|_2 - \|\theta_{rt}^*\|_2) \\
& \geq -\lambda_1 \sum_{t \in N(r)} \sqrt{\nu_{rt}} \|u_{rt}\|_2 \\
& \geq -\lambda_1 d_r \sqrt{\nu_{r \max}} \|u_S\|_{\infty, 2} \\
& = -\frac{5 d_r \nu_{r \max}}{C_{\min}} \lambda_1 (\lambda_1 + \lambda_2)
\end{aligned} \tag{58}$$

Also,

$$\begin{aligned}
& \lambda_2 (\|(\theta_{r\cdot}^*)_S + u_S\|_1 - \|(\theta_{r\cdot}^*)_S\|_1) \\
& \geq -\lambda_2 \|u_S\|_1 \\
& \geq -\lambda_2 \|u_S\|_1 \\
& \geq -\lambda_2 d_r \sqrt{\nu_{r \max}} \|u_S\|_{\infty, 2} \\
& = -\frac{5 \nu_{r \max}}{C_{\min}} d_r \lambda_2 (\lambda_1 + \lambda_2)
\end{aligned} \tag{59}$$

On the other hand, by Taylor's approximation of $\nabla^2 \ell$, there exists $\alpha_{jk} \in [0, 1]$ and \tilde{u}_{jk}^i between $\theta_{\setminus r}$ and $\theta_{\setminus r} + v u_S$ such that

$$\begin{aligned}
& \Lambda_{\min}(\nabla^2 \ell((\theta_{r\cdot}^*)_S + v u_S)) \\
& \geq \min_{\beta \in [0, 1]} \Lambda_{\min}(\nabla^2 \ell((\theta_{r\cdot}^*)_S + \beta u_S)) \\
& \geq \Lambda_{\min}(Q_{SS}^n) - \\
& \max_{v \in [0, 1]} \max_{\|y\| \leq 1} \left\{ \frac{1}{n} \sum_{i=1}^n \sum_{j, k, l, t, h, s, m, s', m'} \alpha_{jk} v (\nabla^3 A_r(\tilde{u}_{jk}^i))_{jkl} \right. \\
& \quad \left. u_{rt;lh} B_{th}(X_t^i) y_{s,m,j} B_{sm}(X_s^i) B_{s'm'}(X_{s'}^i) y_{s',m',j'} \right\}
\end{aligned} \tag{60}$$

Consider the event ε_1 as defined in the previous proof. We know that $P(\varepsilon_1) \geq 1 -$

$c_1 p'^{-3} (\sum_t m_t)$. Conditioned on ε_1 and using Assumption 5 we arrive to the following:

$$\begin{aligned}
& \Lambda_{\min} (\nabla^2 \ell ((\theta_{r.}^*)_S + v u_S)) \\
& \geq C_{\min} - 4 \log p' \|u_S\|_1 D_{\max} \nu_{r \max} k_3(n, p) \\
& \geq C_{\min} - 4 \log p' d_r \sqrt{\nu_{r \max}} \|u_S\|_{\infty, 2} D_{\max} \nu_{r \max} k_3(n, p) \\
& \geq \frac{C_{\min}}{2}
\end{aligned} \tag{61}$$

where the last inequality holds for $\lambda_1 + \lambda_2 \leq \frac{C_{\min}^2}{40 \log p' D_{\max} d_r k_3(n, p) \nu_{r \max}^2}$.

Finally using the above bounds we arrive at the following:

$$F(u_S) \geq d_r \nu_{r \max} \frac{5}{C_{\min}} (\lambda_1 + \lambda_2)^2 \left(-1 - \frac{1}{4} + \frac{5}{2} \right) > 0 \tag{62}$$

Therefore

$$\|(\theta_{r.}^*)_S - (\hat{\theta}_{r.})_S\|_{\infty, 2} \leq \frac{5 \sqrt{\nu_{r \max}}}{C_{\min}} (\lambda_1 + \lambda_2)$$

□

Lemma 4. Suppose that $\lambda_1 + \lambda_2 \leq \frac{\alpha}{2-\alpha} \frac{\sqrt{\nu_{r \min}} C_{\min}^2}{400 \nu_{r \max}^{\frac{5}{2}} \log p' D_{\max} k_3(n, p) d_r}$ and $\|W_{\setminus r}^n\|_{\infty, 2} \leq \frac{\alpha (\lambda_1 + \lambda_2) \sqrt{\nu_{r \min}}}{4(2-\alpha)}$, then,

$$P \left(\frac{\|R^n\|_{\infty, 2}}{\lambda_1 + \lambda_2} \leq \frac{\alpha \sqrt{\nu_{r \min}}}{4(2-\alpha)} \right) \geq 1 - c p'^{-3} \left(\sum_t m_t \right) \tag{63}$$

for some constant $c > 0$.

Proof. Recall that $R_j^n = [\nabla^2 \ell(\theta_{r.}^*; \mathcal{D}) - \nabla^2 \ell(\bar{\theta}_{r.}^j; \mathcal{D})]_j^T (\hat{\theta}_{r.} - \theta_{r.}^*)$ where $[\cdot]_j^T$ denotes the j-th row of a matrix. Let us also refer to $R_{t;jk}^n$ to the coordinate of R^n corresponding to $\theta_{rj;tk}$. Then,

$$\begin{aligned}
R_{t;jk}^n &= \left[\frac{1}{n} \sum_{i=1}^n B_{tk}(X_t^i) \left[\nabla^2 A_r \left(\theta_{r.}^* + \sum_{s \neq r} \theta_{rs}^* B_s(X_s^i) \right) - \right. \right. \\
& \quad \left. \left. \nabla^2 A_r \left(\bar{\theta}_{r.}^{t;jk} + \sum_{s \neq r} \bar{\theta}_{rs}^{t;jk} B_s(X_s^i) \right) \right]_j \otimes B_{r.}^i \right]_j^T (\hat{\theta}_{r.} - \theta_{r.}^*)
\end{aligned} \tag{64}$$

with $B_{r.}^i$ the vector of sufficient statistics evaluate at the i-th sample. Intoroducing the notation $\langle \theta_{r.}, B_{r.}^i \rangle =: \theta_r + \sum_{s \neq r} \theta_{rs} B_s(X_s^i)$, from the mean value theorem we obtain

$$\begin{aligned}
& \nabla^2 A_{r;jl} (\langle \theta_{r.}^*, B_{r.}^i \rangle) - \nabla^2 A_{r;jl} (\langle \bar{\theta}_{r.}^{t;jk}, B_{r.}^i \rangle) = \\
& -v_{j,l}^i [\langle \hat{\theta}_{r.} - \theta_{r.}^*, B_{r.}^i \rangle] (\nabla^3 A_r)_{jl:} (\langle \bar{\theta}_{r.}^{t;jk}, B_{r.}^i \rangle)
\end{aligned} \tag{65}$$

Therefore, combining (64) with (65) and using basic properties of krockner product we

ontain that

$$\begin{aligned} \left| R_{t;jk}^n \right| &\leq \frac{1}{n} \sum_{i=1}^n |B_{tk}(X_t^i)| \nu_{r \max} k_3(n, p) D_{max} \|\hat{\theta}_{r \cdot} - \theta_{r \cdot}^*\|_2^2 \\ &\leq 4 \log p' \nu_{r \max} k_3(n, p) D_{max} \|\hat{\theta}_{r \cdot} - \theta_{r \cdot}^*\|_2^2 \end{aligned} \quad (66)$$

which implies

$$\begin{aligned} \|R_t^n\|_{\infty, 2} &\leq \\ 4 \sqrt{\nu_{r \max}} \log p' \nu_{r \max} k_3(n, p) D_{max} \|\hat{\theta}_{r \cdot} - \theta_{r \cdot}^*\|_2^2 &\leq \\ 4 \sqrt{\nu_{r \max}} \log p' \nu_{r \max} k_3(n, p) D_{max} \frac{25 d_r \nu_{r \max}}{C_{\min}^2} \lambda_1^2 &\leq \\ \frac{\lambda_1 + \lambda_2}{4} \frac{\alpha \sqrt{\nu_{r \min}}}{2 - \alpha} & \end{aligned} \quad (67)$$

with probbility at least $1 - cp'^{-3} (\sum_t m_t)$. \square

We now prove theorem 2 using lemmas 2-4. Recalling that $Q^n = \nabla^2 \ell(\theta_{r \cdot}^*; \mathcal{D})$ and the fact that we have set $(\hat{\theta}_{r \cdot})_{S^c} = 0$ in our primal-dual construction, we can rewrite condition (43) as the following equations:

$$\begin{aligned} Q_{S^c S}^n [(\hat{\theta}_{r \cdot})_S - (\theta_{r \cdot}^*)_S] &= \\ W_{S^c}^n - \lambda_1 \sum_{t \notin N(r)} \sqrt{\nu_{rt}} \hat{Z}_{1,rt} - \lambda_2 \hat{Z}_{2,S^c} + R_{S^c}^n & \end{aligned} \quad (68)$$

$$\begin{aligned} Q_{SS}^n [(\hat{\theta}_{r \cdot})_S - (\theta_{r \cdot}^*)_S] &= \\ W_S^n - \lambda_1 \sum_{t \in N(r)} \sqrt{\nu_{rt}} \hat{Z}_{1,rt} - \lambda_2 \hat{Z}_{2,S} + R_S^n & \end{aligned} \quad (69)$$

Since the matrix Q_{SS}^n is invertible, the conditions (68) and (69) can be rewritten as

:

$$\begin{aligned} Q_{S^c S}^n (Q_{SS}^n)^{-1} [W_S^n - \lambda_1 \sum_{t \in N(r)} \sqrt{\nu_{rt}} \hat{Z}_{1,rt} - \lambda_2 \hat{Z}_{2,S} + R_S^n] &= \\ W_{S^c}^n - \lambda_1 \sum_{t \notin N(r)} \sqrt{\nu_{rt}} \hat{Z}_{1,rt} - \lambda_2 \hat{Z}_{2,S^c} + R_{S^c}^n & \end{aligned} \quad (70)$$

Rearranging yields the following condition:

$$\begin{aligned} \lambda_1 \sum_{t \notin N(r)} \sqrt{\nu_{rt}} \hat{Z}_{1,rt} &= \\ W_{S^c}^n + R_{S^c}^n - Q_{S^c S}^n (Q_{SS}^n)^{-1} [W_S^n + R_S^n] - & \\ \lambda_2 \hat{Z}_{2,S^c} + Q_{S^c S}^n (Q_{SS}^n)^{-1} [\lambda_1 \sum_{t \in N(r)} \sqrt{\nu_{rt}} \hat{Z}_{1,rt} + \lambda_2 \hat{Z}_{2,S}] & \end{aligned} \quad (71)$$

Strict Dual Feasibility: we now show that for the dual sub-vector \hat{Z}_{1,S^c} , we have $\|\hat{Z}_{1,S^c}\|_{\infty, 2} < 1$. We get the following equation from 71, by applying triangle inequality:

$$\begin{aligned}
& \lambda_1 \sqrt{\nu_{r \min}} \left\| \hat{Z}_{1,Sc} \right\|_{\infty,2} \leq \\
& \lambda_1 \left\| \sum_{t \notin N(r)} \sqrt{\nu_{rt}} \hat{Z}_{1,rt} \right\|_{\infty,2} \leq \\
& \left[\|W^n\|_{\infty,2} + \|R^n\|_{\infty,2} \right] \left(1 + \|Q_{ScS}^n (Q_{SS}^n)^{-1}\|_{\infty,2} \sqrt{d_r} \right) \\
& + \lambda_2 \sqrt{\nu_{r \max}} \\
& + \|Q_{ScS}^n (Q_{SS}^n)^{-1}\|_{\infty,2} [(\lambda_1 + \lambda_2) \sqrt{d_r \nu_{r \max}}]
\end{aligned} \tag{72}$$

where $\nu_{r \min} = \min_t \nu_{rt}$, $\nu_{r \max} = \max_t \nu_{rt}$ and $d_r = |N(r)|$ Using mutual incoherence bound 2 on the above equation gives us:

$$\begin{aligned}
& \left\| \hat{Z}_{1,Sc} \right\|_{\infty,2} \leq \\
& \frac{1}{\lambda_1 \sqrt{\nu_{r \min}}} \left[\|W^n\|_{\infty,2} + \|R^n\|_{\infty,2} \right] (2 - \alpha) \\
& + \frac{\lambda_2 \sqrt{\nu_{r \max}}}{\lambda_1 \sqrt{\nu_{r \min}}} \left[1 + \|Q_{ScS}^n (Q_{SS}^n)^{-1}\|_{\infty,2} \sqrt{d_r} \left(\frac{\lambda_1}{\lambda_2} + 1 \right) \right]
\end{aligned} \tag{73}$$

Using the previous lemmas we obtain the following:

$$\begin{aligned}
& \left\| \hat{Z}_{1,Sc} \right\|_{\infty,2} \leq \frac{1}{2\lambda_1} [\alpha (\lambda_1 + \lambda_2)] \\
& + \frac{\lambda_2 \sqrt{\nu_{r \max}}}{\lambda_1 \sqrt{\nu_{r \min}}} \left[1 + \frac{m_{min}}{m_{max}} (1 - \alpha) \left(\frac{\lambda_1}{\lambda_2} + 1 \right) \right]
\end{aligned} \tag{74}$$

If $\lambda_2 < \left(\frac{\alpha}{2 - \alpha + 2 \frac{\sqrt{\nu_{r \max}}}{\sqrt{\nu_{r \min}}}} \right) \lambda_1$, then,

$$\left\| \hat{Z}_{1,Sc} \right\|_{\infty,2} < 1 \tag{75}$$

We have shown that the dual is strictly feasible with high probability and also the solution is unique. And hence based on Lemma 1 the method correctly excludes all edges not in the set of edges.

Correct Neighbourhood Recovery: To show that all correct neighbours are recovered, it suffices to show that

$$\|(\theta_{r^*})_S - (\hat{\theta}_{r^*})_S\|_{\infty,2} \leq \frac{\theta_{min}}{2}$$

where $\theta_{min} = \min_{t \in V \setminus r} \|\theta_{rt}\|_2$.

Using Lemma 3 we can show the above inequality holds if $\theta_{min} \geq \frac{10 \sqrt{\nu_{r \max}}}{C_{min}} (\lambda_1 + \lambda_2)$

C Full MyFitnessPal Graph

Figure 3 shows a high-level view of the entire VS-MRF learned from the MyFitnessPal food database. The three macro-nutrients (fat, carbs, and protein) correspond to the




Figure 3: Full MRF learned from the MyFitnessPal food database. The hubs correspond to point-inflated gamma nutrient nodes, with the three largest hubs being the macro-nutrients (fat, carbs, and protein).

three largest hubs with the remaining nine micro-nutrients representing smaller hubs.

References

- S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2009.
- S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- J. Friedman, T. Hastie, and R. Tibshirani. A note on the group lasso and a sparse group lasso. *arXiv preprint arXiv:1001.0736*, 2010. URL <http://arxiv.org/abs/1001.0736>.
- A. Jalali, P. Ravikumar, V. Vasuki, and S. Sanghavi. On learning discrete graphical models using group-sparse regularization. *AI STAT*, 2011.

- S. L. Lauritzen. *Graphical models*. Oxford University Press, 1996.
- N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2013.
- N. Simon, J. Friedman, T. Hastie, and R. Tibshirani. A sparse-group lasso. *Journal of Computational and Graphical Statistics*, 22(2):231–245, 2013.
- E. B. Sudderth, A. T. Ihler, M. Isard, W. T. Freeman, and A. S. Willsky. Nonparametric belief propagation. *Communications of the ACM*, 53(10):95–103, 2010.
- D. Vats and J. M. Moura. Finding non-overlapping clusters for generalized inference over graphical models. *Signal Processing, IEEE Transactions on*, 60(12):6368–6381, 2012.
- M. J. Wainwright. Sharp thresholds for noisy and high-dimensional recovery of sparsity using ℓ_1 -constrained quadratic programming (lasso). *IEEE Transactions on Information Theory*, pages 2183–2202, 2009.
- M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1-2):1–305, 2008.
- E. Yang, G. Allen, Z. Liu, and P. Ravikumar. Graphical models via generalized linear models. In *Advances in Neural Information Processing Systems*, pages 1358–1366, 2012.
- E. Yang, Y. Baker, P. Ravikumar, G. Allen, and Z. Liu. Mixed graphical models via exponential families. *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, pages 1042–1050, 2014.